UNSUPERVISED LEARNING II: NETWORK THEORY

Gueorgui S. Konstantinov, PhD
Senior Portfolio Manager, Fixed Income and Currencies, DekaBank

Agathe Sadeghi, PhD
Stevens Institute of Technology, School of Business

Introduction

Decision making and risk assessment have traditionally been grounded in core financial models. These models have encountered challenges as global interdependencies have grown in complexity as financial markets evolve. Many classical models approach the market as either a collective system with uniform interactions or groups of unconnected systems where actors make autonomous decisions. These overly simplified approaches have proved incorrect in many incidents, such as the financial crisis of 2007–2008 and the COVID-19 pandemic. The failure of Lehman Brothers and the near-collapse of AIG highlighted how distress is capable of rapidly spreading from one region of the system to another, revealing hidden links and vulnerabilities. Conventional models did not foresee these vulnerabilities. These events revealed the sobering truth that the interconnectivity of financial institutions and the system is much more intricate than the models of the time could capture and that the interexposures between institutions are crucial to the systemic stability of the structure. This crisis proves that models are blind to emerging risks if relationships between players add up in unpredictable asymmetric ways. To respond to the challenge of managing systemic risk, one needs tools that identify and structure relationships to analyze connections, no matter how unconventional the direction might be.[1]

A powerful and intuitive framework for overcoming these constraints is provided by network theory, a discipline with roots in graph theory. It offers mathematical language for investigating systems made up of distinct entities (represented as nodes or vertices) and the connections or interactions among them (represented as links or edges). This method has shown promise in a wide range of fields, including biology, computer science, transportation, and the social sciences.

Network analysis is a natural fit for financial systems. Markets, assets, financial institutions, and even information flows are all interconnected. These relationships create intricate networks that support the dynamics of the market. Practitioners can examine the true structure of market interactions and go beyond oversimplified assumptions by depicting these varied financial relationships as networks. The following are examples of networks:

- Interbank lending networks: A vital network for distributing liquidity is created when banks lend to and borrow from one another.

[1] Fundamental literature on networks includes Wasserman and Faust (1994), Newman (2010), Barabási (2016), and Borgatti, Everett, Johnson, and Agneessens (2022). Financial networks are extensively analyzed by Diebold and Yilmaz (2015).

- Asset and factor networks: By comparing price movements of stocks, bonds, commodities, currencies, or factors, one can identify community structures, relationships, associations, and interactions that improve diversification.
- Ownership networks: Ownership links are created by businesses that own stock in other businesses.
- Derivative networks: Counterparty exposures in derivative contracts, such as credit default swaps (CDSs), create a web of contingent liabilities.
- Bank–firm networks: The financial sector is connected to the actual economy through lending relationships between banks and nonfinancial firms.
- Supply chain networks: Along supply chains, there are dependencies and financial flows.

By demonstrating how information spreads, how sentiment changes, and how various market segments affect one another, network analysis provides a potent lens through which to view market dynamics. It is a crucial tool for evaluating systemic risk, a viewpoint that central banks and regulators have now widely embraced because it enables modeling of contagion pathways and identification of systemically important institutions. It also supports diversification through cluster detection, feature selection, importance score determination, and unveiling hidden relationships between assets. Network models also provide insight into how news and analyst sentiment affect asset prices by tracking the flow of information through markets.

This chapter introduces key network theory concepts and their practical use in finance. It covers community detection to find hidden asset groups, centrality measures to identify influential actors, and fundamental structures such as nodes and edges. It also investigates network dynamics for modeling contagion and systemic risk. Formal definitions are paired with actual financial examples to discuss applications in investment management, such as portfolio construction, market prediction, and pattern recognition.

Concepts

Knowing the fundamentals of graph theory is necessary to comprehend financial systems from a network perspective. These elements offer the framework for modeling complex dependencies.[2]

Nodes

Nodes or vertices are the basic components or entities that make up a network. They stand in for the actors or objects in the system that is being modeled. The definition and selection of a node are important decisions that depend solely on the particular financial system under study. Numerous entities can be represented by nodes:

- Institutions, such as central banks, investment banks, mutual funds, and hedge funds
- Corporations, such as nonfinancial businesses that participate in supply chains or credit relationships
- Assets and factors, such as individual stocks, bonds, currencies, commodities, or derivatives, such as CDSs

[2] A detailed explanation of these metrics can be found in Newman (2010) and Barabási (2016).

- People, such as market traders, stock analysts, board members, firm employees, and even politicians in policy networks
- Geographic/political entities, such as nations, jurisdictions, communities, or areas involved in global networks of capital flows or trade

Financial network nodes are more than just abstract points; they also carry characteristics that give analysis crucial context. These attributes, such as the total assets of a bank, the sector or volatility of a stock, the credit rating of a company, or the GDP of a nation, aid in differentiating nodes and influence how network relationships are interpreted. A key first step in creating meaningful financial network models is defining nodes and their relevant attributes because this process establishes what relationships can be captured and what insights can be obtained.

Edges

The connections, relationships, interactions, or dependencies between pairs of nodes in a network are represented by edges, also referred to as links or arcs. They represent the connections between the nodes. In finance, they capture the way financial entities interact with one another. The following are examples:

- Lending and borrowing: In an interbank network, an edge can stand in for a loan from Bank A to Bank B or a bank credit line to a business.
- Association: Two stocks whose returns show a high degree of association can be connected by an edge. Association metrics, such as pairwise correlations between asset classes, are the most widely used tool to measure association and to draw edges in financial networks.
- Counterparty: In a derivative contract (such as a CDS), an edge can stand in for the possible loss exposure between two parties.
- Informational: Two stocks may be linked if they are covered by the same financial analyst.
- Ownership: An edge may indicate that Firm A owns a sizable portion of Firm B's stock.
- Affiliation: Edges can connect businesses that share board members.

Edges in financial networks can have attributes that are essential for network analysis. The difference between directed and undirected edges is a crucial one. A relationship with a distinct origin and destination is indicated by directed edges, such as when Bank A lends to Bank B. The relationship's direction is vital in these situations, particularly when simulating influence or contagion. Undirected edges, in contrast, show symmetrical or reciprocal relationships in which order is irrelevant, such as the correlation between two stocks.

Additionally, edges can be weighted or unweighted. Weighted edges, such as trading volume, correlation coefficients, or loan size, show how strong a relationship is. Ignoring these weights can result in a substantial loss of information because they provide important detail. Conversely, unweighted edges do not quantify the strength of a relationship; they show only whether one exists. They are helpful in situations where a connection's existence is more significant than its size. The adjacency matrix of a network is the one that carries the information about the pairwise weight structure.[3]

[3] See the next section for detailed information about the adjacency matrix.

To create precise and instructive financial network models, edges must be carefully defined, including their directionality and weighting. The outcomes of network analysis, including contagion simulations and centrality calculations, are directly affected by these characteristics.

Exhibit 1 provides a simple example of the correlation-based network of the Dow Jones Industrial Average Index. The network consists of 30 nodes (stocks in the index), and according to the simple algorithm, there are 570 links between the nodes. Given the 30 nodes in the networks, the total number of possible links between all 30 nodes is 30 × 29 = 870. In other words, the algorithm's output of 570 links between the nodes shows that not all nodes share a pairwise link between them. For example, American Express might not have a link to Amazon. Similarly, the edge between, say, Boeing and Visa might not exist.

Exhibit 1. A Simple Correlation-Based Network of the Dow Jones Industrial Average Index

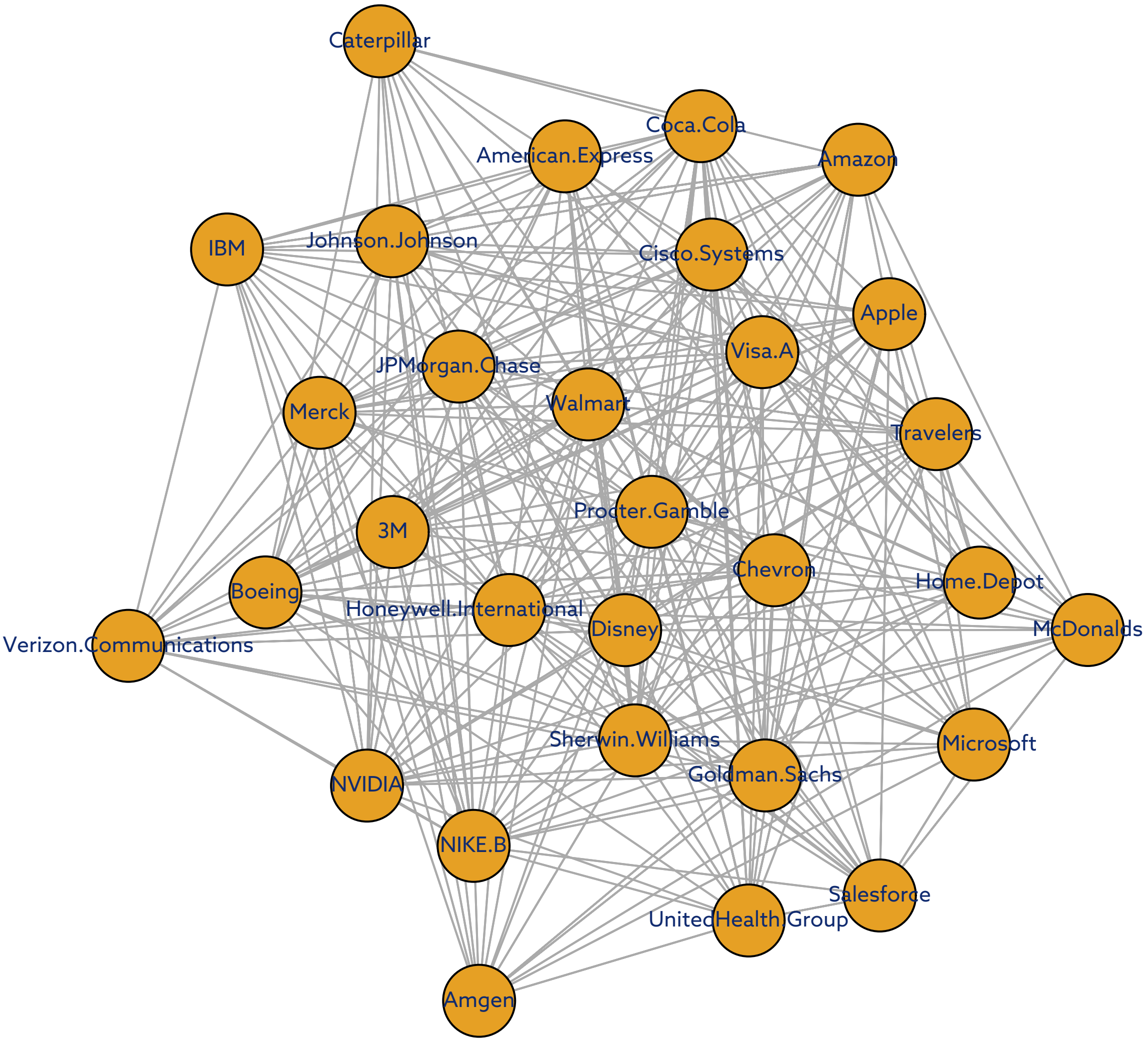

Source: Bloomberg, LLC.

Path, Trails, Walks, Geodesic Distances, and Network Structure

Higher-level insights can be gained from the general arrangement and connectivity patterns of a network, which go beyond individual nodes and edges. A crucial distinction lies in the definitions of paths, trails, walks, and geodesic distances.

A path is a route in a network consisting of a sequence of nodes connected by edges. The terms *geodesic*, *trail*, *path*, and *walk* describe different types of node–edge sequences in a network. They differ in their constraints and the connectivity characteristics they represent.

A geodesic is the shortest path between two nodes in a network. It represents the most direct route, taking into account the network's structure and any weights assigned to its edges (if applicable). Geodesic distances are used to calculate various network metrics, such as average path length, network diameter, and centrality measures. The length of a path is determined by the number of edges it traverses; a path with a length of 1 consists of a single edge. For centrality metrics, such as closeness and betweenness, the shortest path between nodes is critical because it minimizes either total weight (in weighted graphs) or the number of edges (in unweighted graphs).

Paths are useful in finance for modeling complex transaction sequences, analyzing the propagation of shocks through counterparties, and tracking the flow of information from analysts to investors. Konstantinov, Aldridge, and Kazemi (2023) extensively discussed the flow properties of financial networks. A walk is a sequence of nodes and edges in which each edge is adjacent to the previous one and both nodes and edges can be revisited. Walks can include loops and repeated edges. This property is typical in financial networks. A trail is similar to a walk, but no edge may be repeated, although nodes can still be revisited. For instance, a stock can be bought multiple times from the same brokerage house, but the exact transaction (edge) cannot be repeated. A path further restricts the sequence such that no node or edge is repeated. This type of sequence can be quite limiting in financial networks.

According to Newman (2010), some directed networks have a special structure called cycles. A cycle is a walk in which no node is repeated (except for the starting/ending node), forming a closed loop where all edges point in the same direction. Networks without such loops are referred to as *acyclic directed networks*, which are widely used in causal inference and prediction.

The primary differences between these sequences are the rules about revisiting nodes and edges. When analyzing flow processes in networks, each type of sequence can represent different aspects of connectivity, traversal, or the spread of information. For example, geodesics can indicate the most efficient information flow between nodes. Thus, applying geodesic metrics to financial networks provides a generalizable framework for modeling relationships. The simplest way to model such relationships is through pairwise correlation metrics.

An important consideration in flow processes is whether money, information, or services reach a node only once, simultaneously, or multiple times. Financial flows often arrive at nodes repeatedly over time. For example, a portfolio's exposure to certain assets may change frequently depending on the economic environment. In this context, *systematic* financial risk can become *systemic*, simultaneously impacting many nodes via multiple links, not necessarily along a single geodesic path.

Similarly, information can reach multiple points at once, spreading by replication rather than linear transfer and influencing many nodes simultaneously.

In contrast, goods and packages are often transferred along geodesic paths—a model that may not always be appropriate for finance, where assets, factors, or institutions may be disconnected. Consider a hedge fund using a sophisticated quantitative strategy with minimal market exposure. In a hedge fund network, this fund might not be linked to others based on correlation coefficients, meaning no directed path exists. Directionality in networks implies causality, a key concept in financial markets, although it is often difficult to detect. The direction of money, information, and technological impacts (e.g., high-frequency trading) reflects this causality. The processes of paths, trails, and walks help define how relationships form and evolve between nodes.

The underlying flow process—whether involving information, rumors, money, goods, services, orders, transactions, or financial contagion—is critical for network modeling. A major distinction lies in whether flows follow geodesics, trails, paths, or walks. In financial markets, money typically follows walks rather than trails, which is a defining characteristic of financial flow and essential for network analysis.

Based on mathematical graph theory, a network (or graph) is a collection of nodes (vertices) connected by edges (links). A network G is defined as a set of nodes n and edges m, formally expressed as $G = (n, m)$.

There are two primary ways to construct a financial network:

- by modeling the data at the edge level, keeping the nodes fixed, or
- by modeling the network at the node level, with edges defined between them.

Mathematically, relationships between nodes are often represented using a matrix. Financial networks can be modeled by an adjacency matrix, where rows and columns represent nodes and each cell indicates the presence or strength of a connection.

In unweighted networks, entries in the adjacency matrix are binary (0 or 1), whereas in weighted networks, these entries are real numbers reflecting the strength or weight of the connection. Generally, in an unweighted network, the adjacency matrix has cell values $[i, j]$ equal to 1 if there is a link from node i to node j and 0 otherwise. Using this notation, the result is a simple network whose main diagonal reflects self-edges. An example of self-edge (main diagonal of nonzeros) indicates that the node has a self-connectedness. Mathematically for a graph, $G = (n, m)$, where n is the number of nodes and m is the number of edges, the $n \times n$ adjacency matrix is denoted by $[A]_{ij}$ and is estimated using following rule:

$$a_{ij} = \begin{cases} 1 \text{ if node } i \text{ and } j \text{ are connected, or } \{i, j\} \in m \\ \qquad\qquad 0 \text{ otherwise} \end{cases}$$

The adjacency matrix for an undirected network consisting of 30 nodes for the Dow Jones Industrial Index is shown in **Exhibit 2**.

Konstantinov and Fabozzi (2025) provided an extensive review of the useful and practical ways to construct financial networks. The approaches used to estimate the links between the assets include correlation coefficients, regression models, econometric models, probabilistic models that derive probabilities, evaluation of transaction and money flow volumes, assets under management, and all possible information that applies to exchange in financial markets.

Exhibit 2. Adjacency Matrix

Goldman.Sachs	0 0 0 0 0 0 1 1 1 0 1 1 1 0 1 0 1 1 1 0 1 1 0 1 0 0 1 1 1 1
UnitedHealth.Group	0 0 0 0 0 1 1 1 1 0 1 1 1 0 1 0 1 1 1 0 1 1 0 1 0 0 1 1 1 0
Microsoft	0 0 0 0 0 1 1 1 1 0 1 1 1 0 1 0 1 0 1 0 1 1 0 1 0 1 1 1 1 0
Home.Depot	0 0 0 0 0 1 0 1 1 0 1 1 1 0 1 0 1 0 1 0 1 1 0 1 0 1 1 1 1 0
Caterpillar	0 0 0 0 0 1 0 1 0 0 0 0 0 0 1 0 1 1 1 1 1 1 0 0 0 1 1 1 1 0
Sherwin.Williams	0 1 1 1 1 0 1 1 1 1 1 1 1 1 0 0 1 1 1 0 1 1 0 1 1 1 1 1 1 1
Salesforce	1 1 1 0 0 1 0 0 0 1 0 1 0 1 1 0 0 0 1 1 1 0 1 0 1 1 1 0 1 0
Visa.A	1 1 1 1 1 1 0 0 0 1 0 0 1 1 1 1 0 1 1 1 0 0 1 1 1 1 0 1 0 0
American.Express	1 1 1 1 0 1 0 0 0 0 0 0 1 1 1 1 0 1 1 1 0 0 1 1 1 1 0 1 0 0
McDonalds	0 0 0 0 0 1 1 1 0 0 1 0 1 0 1 0 1 0 1 0 1 1 0 1 0 1 1 1 1 0
Amgen	1 1 1 1 0 1 0 0 0 1 0 0 0 1 1 1 0 0 1 1 0 0 1 1 1 1 0 0 0 0
Apple	1 1 1 1 0 1 1 0 0 0 0 0 0 1 1 1 0 1 1 1 0 0 1 1 1 1 0 1 0 0
Travelers	1 1 1 1 0 1 0 1 1 1 0 0 0 1 1 0 0 0 1 1 1 1 1 1 1 1 0 1 1 0
JPMorgan.Chase	0 0 0 0 0 1 1 1 1 0 1 1 1 0 1 0 1 1 1 0 1 1 0 1 1 0 1 1 1 1
Honeywell.International	1 1 1 1 1 0 1 1 1 1 1 1 1 1 0 1 1 1 0 1 1 1 1 1 1 1 1 0 1 1
IBM	0 0 0 0 0 0 0 1 1 0 1 1 0 0 1 0 1 0 1 1 0 1 0 1 1 1 1 1 1 1
Amazon	1 1 1 1 1 1 0 0 0 1 0 0 0 1 1 1 0 0 1 1 1 0 1 1 1 1 0 0 0 0
Boeing	1 1 0 0 1 1 0 1 1 0 0 1 0 1 1 0 0 0 1 1 1 0 1 1 1 1 1 0 1 0
Procter.Gamble	1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 1 1 1 0 0 1 1 1 1 1 0 1 1 1 1
Johnson.Johnson	0 0 0 0 1 0 1 1 1 0 1 1 1 0 1 1 1 1 0 0 1 1 0 1 1 0 1 1 1 1
Chevron	1 1 1 1 1 1 1 0 0 1 0 0 1 1 1 0 1 1 1 1 0 0 1 1 1 1 1 1 0 1
NVIDIA	1 1 1 1 1 1 0 0 0 1 0 0 1 1 1 1 0 0 1 1 0 0 1 1 1 1 1 0 0 1
3M	0 0 0 0 0 0 1 1 1 0 1 1 1 0 1 0 1 1 1 0 1 1 0 0 0 1 1 1 1 1
Disney	1 1 1 1 0 1 0 1 1 1 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 1 1 1 1 1
Merck	0 0 0 0 0 1 1 1 1 0 1 1 1 1 1 1 1 1 1 1 1 1 0 0 0 1 1 1 1 1
Walmart	0 0 1 1 1 1 1 1 1 1 1 1 1 0 1 1 1 1 0 0 1 1 1 1 1 0 1 1 1 1
NIKE.B	1 1 1 1 1 1 1 0 0 1 0 0 0 1 1 1 0 1 1 1 1 1 1 1 1 1 0 0 0 1
Coca.Cola	1 1 1 1 1 1 0 1 1 1 0 1 1 1 0 1 0 0 1 1 1 0 1 1 1 1 0 0 1 1
Cisco.Systems	1 1 1 1 1 1 1 0 0 1 0 0 1 1 1 1 0 1 1 1 0 0 1 1 1 1 0 1 0 0
Verizon.Communications	1 0 0 0 0 1 0 0 0 0 0 0 0 1 1 1 0 0 1 1 1 1 1 1 1 1 1 1 0 0

Some structural features are frequently seen in real-world financial networks:

- A core–periphery structure displays a densely connected core (mainly large players) surrounded by a sparsely connected periphery (relatively smaller players). For instance, large international banks with extensive cross-border ties typically form the core of the global banking system, with national or regional banks that have fewer direct international ties encircling the core.
- Scale-free networks have a power-law degree distribution, which means that although most nodes have few connections, a select few (hubs) have a lot of connections. These hubs frequently have a significant impact on the stability and operation of networks. A few big, highly linked banks serve as hubs in the interbank lending market, which frequently has a scale-free structure. Such a hub's failure may have systemic repercussions that are disproportionately severe.
- Small-world networks are characterized by short average path lengths between any two nodes and high clustering. Fast transmission throughout the network is made possible by this structure. Financial markets might frequently act like small-world networks, enabling information, sentiment, or shocks to spread remarkably quickly.

Understanding these structural patterns aids in understanding the general behavior, resilience, and shock susceptibility of the network. The particular topology has a big impact on how information moves, how fast contagion spreads, and where systemic risks could be concentrated.

How to evaluate networks is an important question, and a graph theoretic toolkit allows us to measure the specific properties of a network. A classical financial network has specific properties, and the average degree, density, centrality scores, reciprocity, and clustering coefficient are the most important. The average degree k_i of a node *i* and k_j of node *j* at time *t* is the most basic structural property representation of the connections of a node in a network:

$$k_{i,t} = k_{j,t} = \sum_i a_{ij,t} = \sum_j a_{ji,t} = \frac{m}{n}.$$

The *density*, or *completeness*, of a network, ς, is given by the ratio between the numbers of current links relative to the number of all possible links between the nodes in the network:

$$\varsigma_t = \frac{m_t}{n_t\,(n_t - 1)}.$$

Here, *m* represents the number of links between nodes in the network and *n* represents the number of nodes. The density values range from 0 (no ties are present) to 1 (a completed network). A density value of 1 means all possible links are used.

The relation between the pairs of nodes in a directed network (a network with direction or flow, indicated by arrows) at time *t* is given by the value of *reciprocity*, ϱ_t. The reciprocity measures whether a node *i* is linked to a node *j* and if a node *j* is also linked to a node *i*. The values range between 0 and 1 (with 1 for fully reciprocal network). Given the entries in the adjacency matrices, the reciprocity is given by

$$\varrho_t = \frac{\sum_{ij} a_{ij,t}\ a_{ji,t}}{m_t}.$$

A node pair (i, j) is called *reciprocal* if ties exist between both nodes in both directions.

Centrality Measures

After the network is defined, finding the most important nodes is a crucial task. Although influence can take many different forms, centrality measures assign scores based on a node's position and connections. Selecting the appropriate measure depends on the information flow process because different measures quantify different aspects of the nodes. We discuss the most important measures of network properties.

Degree Centrality

The simplest measure, degree centrality, counts the number of direct connections a node has. It is denoted as deg(v), where v is the node. A high degree measure denotes a popular or active entity—for instance, a bank with many trading relationships or a frequently traded stock. To compare networks of varying sizes, a normalized version (by number of nodes) is frequently used.

If the network is directed, the number of incoming edges is called *in-degree*, and the number of outgoing edges is referred to as *out-degree*. A node with a high in-degree is one that receives a lot of connections or attention, such as a bank that borrows a lot and is therefore at risk of lender distress or a stock with a lot of buy recommendations. In contrast, a node with a high out-degree is a source of influence or risk, such as a bank lending to numerous people, which could spread distress if it fails.

Degree is often used as a measure of a node's connectedness in networks. In real-world networks, the average degree typically exceeds 1, indicating that a node is connected to more than one other node. A higher average degree or greater overall connectedness increases the likelihood of forming communities within the network.

One limitation of degree centrality is its exclusive focus on direct, immediate connections. This approach does not account for a node's position within the broader network or its indirect influence via the connections of its neighbors. A node may have a high degree but be linked only to minor, peripheral nodes, limiting its actual influence. Furthermore, standard degree centrality ignores the weights of the edges, although this shortcoming can be addressed by defining a weighted degree centrality.

As noted by Konstantinov and Fabozzi (2025), the key advantage of degree centrality—and metrics that assess the degree of all nodes in a network—is its ability to measure the immediate impact of risk at the node level. In contrast, eigenvector centrality captures both direct and indirect long-term influence across the network.

Importantly, many mathematical models in network science place a central focus on degree distribution, which plays a foundational role in understanding the structure and dynamics of complex networks.

Betweenness Centrality

The extent to which a node is situated on the shortest routes between any other pair of nodes in the network is measured by betweenness centrality. Developed by Linton C. Freeman, it gauges a node's function as a bridge or middleman. The fraction of shortest paths between each pair of nodes (u,v) that pass through node s is added up using the following formula:

$$C(b) = \sum_{u \neq v \neq s} \frac{\sigma_{u,v}(s)}{\sigma_{u,v}},$$

where $\sigma_{u,v}(s)$ is the number of shortest paths that go through s and $\sigma_{u,v}$ is the total number of shortest paths between u and v.

The application of betweenness centrality in financial networks depends on the underlying flow processes within the network. Unlike such metrics as closeness centrality, which assume random path selection, betweenness centrality is particularly appropriate for networks in which the flow of resources or information follows specific, predetermined paths.

A key assumption when applying betweenness centrality is that transactions or flows tend to follow the shortest paths between nodes. This assumption does not always hold in financial markets, however, where flows may follow more complex, indirect routes. As such, the application of closeness centrality in financial contexts requires careful consideration and may not be universally appropriate.

Nodes with high betweenness centrality are identified as critical intermediaries that control or facilitate the movement of capital, information, or risk across the network. These may include banks that bridge market segments, clearinghouses, or major dealers in the financial industry. Betweenness centrality can also highlight bottlenecks or key links between otherwise disconnected institutions or investor groups.

A node with low betweenness is typically not essential for linking other nodes in the network. Although high betweenness scores often indicate influence, they may also reflect peripheral connectors between clusters. Therefore, interpreting these scores depends heavily on the context and structure of the specific financial network being analyzed.

Closeness Centrality

Closeness centrality measures the average distance between a node and every other reachable node in the network. This metric, based on the foundational work of Linton C. Freeman, is grounded in the idea that a node's importance is inversely related to its geodesic distance from other nodes. It is calculated as the reciprocal of the sum of the shortest path distances from the node to all other nodes in the network.

A node with a higher closeness score (i.e., a lower average distance) can reach other nodes more quickly, indicating greater efficiency in communication or influence across the network.

In the context of financial networks, assets or institutions with high closeness centrality are assumed to respond more rapidly to information, shocks, or risk events compared with those that have lower scores. For example, a financial institution that is well positioned to quickly disseminate or receive contagion would likely exhibit a high closeness centrality value.

The conventional formula for closeness is the inverse of the sum of distances:

$$C(v) = \frac{N-1}{\sum_{u \neq u'} d(u,u')},$$

where N is the number of nodes and $d(u,u')$ is the shortest path distance between the two nodes.

Differentiation can be challenging in highly connected networks because many nodes may have similar, high closeness scores. This measure is also sensitive to the overall structure of the network. Finding all-pairs shortest paths is necessary for calculations, and for very large networks, this process can be computationally demanding.

Eigenvector Centrality

The idea behind eigenvector centrality is that a node's significance is determined by the significance of the nodes it is connected to, not just by the number of connections it has. A node's eigenvector centrality will be high if it is connected to numerous highly central nodes. The components of the principal eigenvector of the network's adjacency matrix (**A**), which corresponds to the largest eigenvalue (λ), are mathematically known as the centrality scores (x):

$$\mathbf{A}x = \lambda x.$$

The key to identifying systemically significant institutions lies in eigenvector centrality, which captures nodes that are influential not only through their direct connections but also through links to other highly connected and influential nodes. Conversely, nodes with low eigenvector centrality are typically connected only to peripheral or less significant nodes.

An important property of eigenvector centrality is that it is only appropriate for analyzing nodes that are connected to at least one other node. If a node is completely disconnected from the rest of the network, eigenvector centrality assigns it a score of zero. This scoring may not accurately reflect the node's importance, however, because it ignores exogenous factors that can affect a node's role within the broader system.

A simple example illustrates this limitation. Consider an asset such as the risk-free rate, which may be temporarily disconnected from other assets in a financial network. In this case, eigenvector centrality would assign it a score of zero, implying no influence. This implication is misleading, however, because the risk-free rate is determined by exogenous factors—for instance, central bank policy—that have significant system-wide effects and influence the connectivity and behavior of other assets in the network. Thus, although eigenvector centrality is a powerful tool for identifying influential nodes based on endogenous network structure, it should be interpreted carefully, especially in contexts where external forces play a key role in shaping network dynamics.

A direct comparison of four widely used centrality metrics provides insight into how importance scores vary among algorithms. Each centrality measure captures a distinct aspect of a node's role within the network, resulting in different scores for the same node. In the corresponding network visualizations in **Exhibit 3**, larger node sizes represent higher importance scores based on the respective centrality metric. **Exhibit 4** presents the numerical values associated with each node under the four centrality measures. Obviously, a node's centrality depends on the specific centrality metric used. For instance, a node with a high betweenness centrality score may not necessarily have a high eigenvector centrality score, because each metric captures different aspects of a node's role within the network.

Exhibit 3. Overview of the Betweenness, Closeness, Eigenvector, and Degree Centrality of a Dow Jones Industrial Average Correlation-Based Network

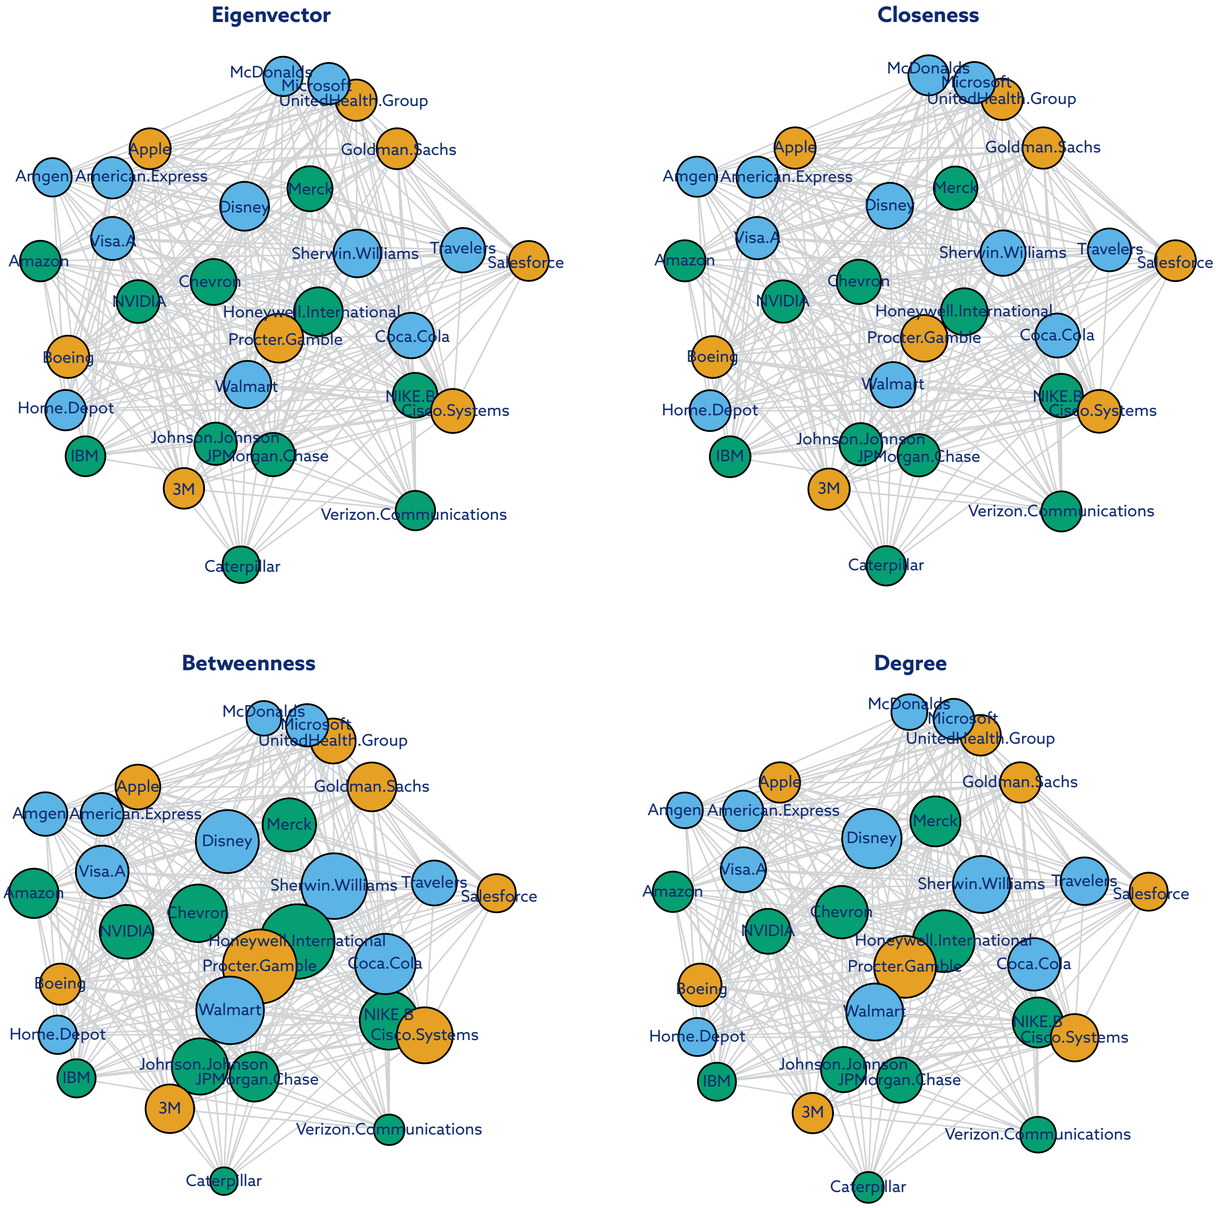

Source: Bloomberg, LLC.

Exhibit 4. Centrality Scores for Betweenness, Closeness, Eigenvector, and Degree Centrality of a Dow Jones Industrial Average Correlation-Based Network

	Eigenvector	Closeness	Betweenness	Degree
Goldman.Sachs	0.69	0.71	0.01	34
UnitedHealth.Group	0.71	0.71	0.01	34
Microsoft	0.72	0.71	0.01	34
Home.Depot	0.69	0.69	0.01	32
Caterpillar	0.56	0.64	0.00	26
Sherwin.Williams	0.95	0.85	0.02	48
Salesforce	0.66	0.69	0.01	32
Visa.A	0.77	0.74	0.01	38
American.Express	0.71	0.71	0.01	34
McDonalds	0.65	0.67	0.01	30
Amgen	0.62	0.67	0.01	30
Apple	0.70	0.71	0.01	34
Travelers	0.83	0.76	0.01	40
JPMorgan.Chase	0.79	0.74	0.01	38
Honeywell.International	0.99	0.91	0.03	52
IBM	0.66	0.69	0.01	32
Amazon	0.69	0.71	0.01	34
Boeing	0.75	0.72	0.01	36
Procter.Gamble	1.00	0.91	0.03	52
Johnson.Johnson	0.75	0.74	0.01	38
Chevron	0.88	0.81	0.02	44
NVIDIA	0.77	0.74	0.01	38
3M	0.69	0.71	0.01	34
Disney	1.00	0.88	0.02	50
Merck	0.85	0.78	0.01	42
Walmart	0.94	0.85	0.02	48
NIKE.B	0.84	0.78	0.02	42
Coca.Cola	0.87	0.81	0.02	44
Cisco.Systems	0.81	0.76	0.01	40
Verizon.Communications	0.65	0.67	0.00	30

A simple comparison illustrates the point. Consider the eigenvector centrality of Goldman Sachs (GS), and compare it with the degree and closeness centrality that GS has in the network. Recall that eigenvector centrality assigns a higher score to a node that is itself connected to other highly connected nodes. Degree centrality measures the total number of links that a specific node has to other nodes. In this respect, GS has a relatively high eigenvector centrality (0.69) and is highly connected to other highly connected nodes, but its degree of 34 is in the lower percentile of degree connectedness. That is, GS has 34 edges to other nodes. The eigenvector centrality, however, indicates that GS is not connected to highly connected nodes.

Community Detection

Network analysis offers powerful tools not only for identifying individually significant nodes but also for uncovering structural patterns, such as clusters of nodes that are more closely connected to each other than to the rest of the network. These groups are commonly referred to as modules, communities, or components. Identifying such communities helps reveal the underlying structure of a network, providing insights into shared characteristics, functional relationships, and potential vulnerabilities within financial systems.

A community is typically defined as a subset of nodes within a network where the density of internal connections is significantly higher than the density of connections between that subset and the rest of the network. The emergence of community structure depends on the network regime. A network is said to be in a connected regime when it forms cohesive components or communities. The average degree of nodes in the network plays a critical role in this dynamic: Once it exceeds a certain threshold, the network begins to form a connected component.

To identify these communities, community detection algorithms aim to maximize the number of intracommunity links while minimizing intercommunity links. These algorithms divide the network into distinct (and sometimes overlapping) communities, and at a coarser resolution, this process effectively reveals the network's global structural organization.

Finding communities in financial networks enables the discovery of significant clusters that may not be visible through node-level analysis alone. The following list provides examples:

- Banking groups: Groups of banks with comparable balance sheets or a high level of interbank activity indicate shared vulnerabilities or concentrated risk.
- Asset communities (clusters): Diversification is aided by identifying sectors or risk exposures by classifying assets based on strong return correlations. Unlike static classifications, such as GICS, these clusters can capture dynamic market structures.
- Functional modules: Communities in supply chains or financial systems frequently represent groups carrying out related tasks.
- Shared interest groups: Communities in financial information networks identify groups that share information sources, tactics, or interests.

Exhibit 5 provides an example for the community detection of the Dow Jones Industrial Average Index according to a correlation-based network and the Cluster Spinglass algorithm according to the highest modularity score.

Exhibit 5. A Correlation-Based Network for the Dow Jones Industrial Average Index

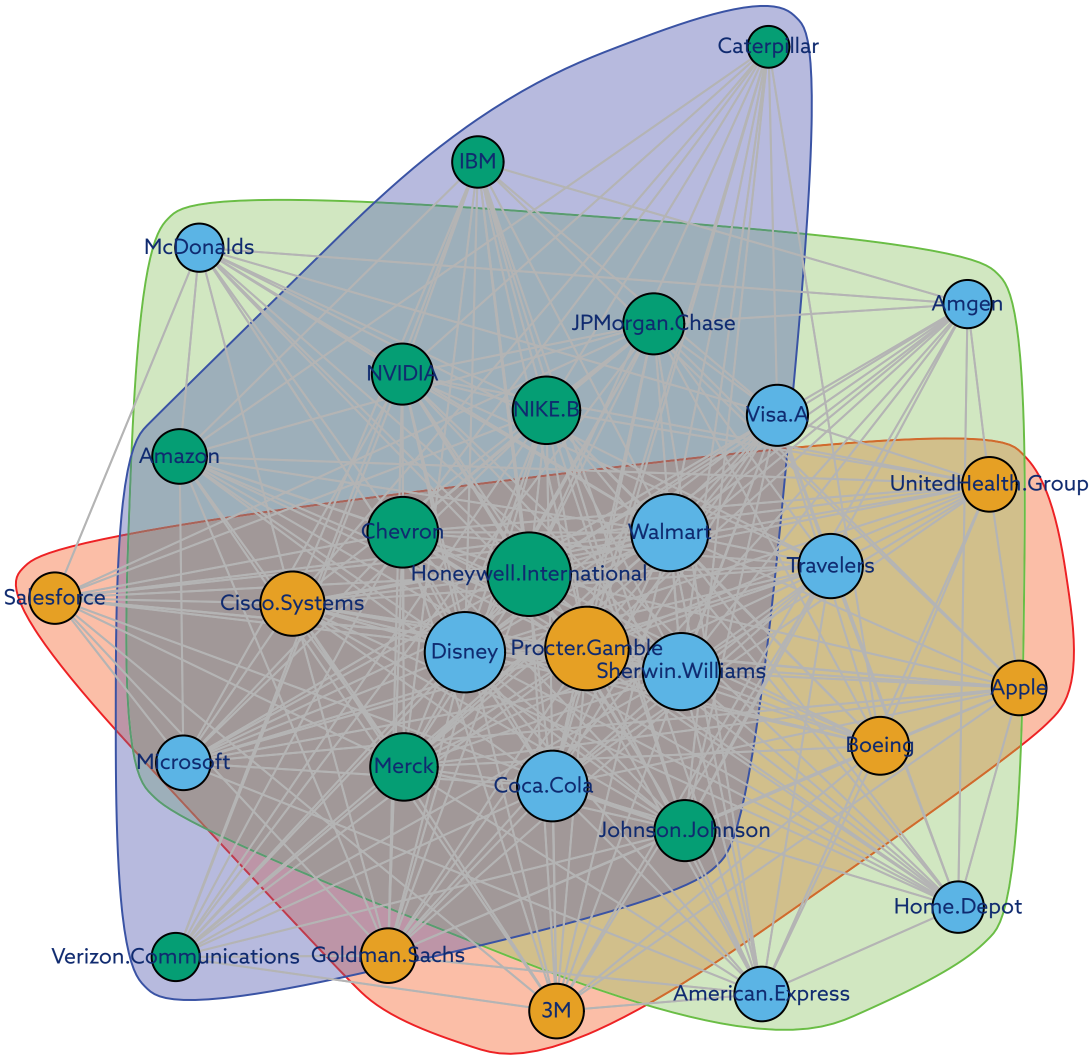

Source: Bloomberg, LLC.

Measuring Community Structure

Assessing the quality of a particular network partition is a major challenge in community detection. The most popular metric for this assessment is modularity.

By comparing the number of edges that actually fall within the suggested communities with the number that would be predicted if edges were dispersed randomly throughout the network while maintaining each node's degree, modularity quantifies how well a network is divided into communities. Strong community structure is indicated by a high modularity score, which means that the identified communities are much more densely connected internally and sparsely connected externally than would be predicted by chance. Usually, modularity values fall between −0.5 and 1. In real-world networks, values higher than roughly 0.3 are frequently regarded as suggestive of substantial community structure.

The following formula suggested by Clauset, Newman, and Moore (2004) is used to determine the modularity, Q, for a network's partition C with adjacency matrix A:[4]

$$Q = \frac{1}{2M} \sum_{i,j} \left[A_{ij} - \frac{k_i k_j}{2M} \right] \delta(C_i, C_j),$$

where A_{ij} is the weight of the edge connecting nodes i and j (1 if connected and unweighted and 0 otherwise). $k_i = \sum_j A_{ij}$ is the weighted degree of node i. $M = \frac{1}{2} \sum_{i,j} A_{ij}$ is the network's total number of edges (or weight sum). The community that node i is assigned to is called C_i. $\delta(C_i, C_j)$ is the Kronecker delta, which is 0 if nodes i and j are not in the same community and 1 otherwise. In this formula, the difference between the expected and actual edge weights for every pair of nodes in the same community is added up.

For many community detection algorithms, modularity serves as both an objective function and an evaluation metric. One widely used algorithm, the Louvain method, specifically seeks the network partition that maximizes the modularity score. The choice and performance of community detection algorithms depend significantly on the type of graph (e.g., directed vs. undirected) and its characteristics.

According to Yang, Algesheimer, and Tessone (2016), algorithm selection should consider network size (number of nodes), number of edges, and the mixing parameter—a measure of how well defined the communities are. Networks with fewer than 1,000 nodes are generally considered small, and most standard algorithms, including Louvain, Spinglass, and fast greedy, are well suited for these. For larger networks, such algorithms as Multilevel, Walktrap, and Infomap may be more appropriate, particularly as computational efficiency becomes critical.

Community Detection Models

A variety of algorithms have been developed to identify communities in networks. These can be broadly categorized by their methodology:

- Bottom-up approaches begin with each node as its own community and merge communities based on such criteria as modularity gain.
- Top-down approaches start with the full network and iteratively split it, often by removing edges with high betweenness centrality.
- Optimization-based methods aim to maximize a quality function (e.g., modularity) using heuristics.

[4] See Newman (2010) and Bech and Atalay (2010).

Other techniques include random walks, spectral clustering, and statistical inference. Commonly used algorithms in financial applications include Louvain, Girvan–Newman, hierarchical clustering, and *k*-means.

Louvain Method

The Louvain algorithm is a greedy optimization technique that operates in two iterative stages to maximize modularity:

- Local optimization phase: Each node begins in its own community. Nodes are moved to neighboring communities if such moves increase modularity. This process continues until no further improvement is possible.
- Aggregation phase: Identified communities are collapsed into "super-nodes," forming a new network with weighted edges reflecting the total connections between groups. The two-phase process repeats on the new network until modularity no longer increases.

The Louvain method is widely applied in finance to group large sets of stocks based on correlation matrices, identifying risk groups or sectors for portfolio construction.

Girvan–Newman Algorithm

This algorithm identifies communities by progressively removing edges with the highest edge betweenness centrality—those that frequently appear in the shortest paths between nodes. In finance, this approach can be useful for detecting correlated asset clusters or uncovering cohesive subgroups within interbank networks.

Hierarchical Clustering

Hierarchical clustering creates a tree-like structure (dendrogram) of clusters. Variants include the following:

- Single linkage—based on the minimum distance
- Average linkage—based on the average distance
- Complete linkage—based on the maximum distance

Typically, correlation is used as the similarity metric. This method is popular for asset clustering.

k-Means Clustering

In *k*-means, nodes are grouped into a predefined number (*k*) of clusters by minimizing the distance between each node and its assigned cluster center. Although simple, this process requires prior specification of *k* and is commonly used for clustering stocks or clients in the financial sector.

The optimal choice of algorithm depends on the network size, network type (weighted/unweighted, directed/undirected), computational resources, and analytical objectives. Because different algorithms may produce distinct community structures from the same dataset, it is crucial to understand their assumptions and test the robustness of the results.

Case Study: Asset Clustering

A well-known application of community detection in finance is grouping financial assets (usually stocks) for the purpose of portfolio diversification. The central idea is that assets in different clusters behave more independently, whereas those in the same cluster exhibit similar behavior (e.g., high return correlation). This strategy can be implemented as follows:

1. Compute pairwise correlations between asset returns over a specified period.
2. Convert the correlation coefficients (ρ) into a distance metric (**d**) using a transformation such as $\mathbf{d}_{ij} = \sqrt{2(1-\rho_{ij})}$. This step is necessary because correlation coefficients should be converted to Euclidean distance metrics.
3. Apply a community detection algorithm (e.g., Louvain, Spinglass, fast greedy, leading eigenvector, Walktrap, or Multilevel) on the resulting network.
4. Select representative assets from each cluster. Selection criteria may include node centrality, cluster size, or other financial metrics.

The effectiveness of a cluster-based investment strategy can be assessed by comparing portfolio performance with benchmarks or peer strategies. Standard evaluation tools include the Sharpe ratio, the information ratio, and factor exposure analysis. These metrics help determine whether the identified clusters and community structures contribute meaningfully to risk-adjusted returns.

Network Dynamics

Financial systems are inherently dynamic and evolving. Modeling these dynamics—especially the spread of shocks and the emergence of systemic risk—is a core application of network theory. Mathematical models from network analysis provide essential tools for evaluating these network dynamics across different market cycles. Key structural metrics used to capture and assess financial network behavior include the following:

- Edge density
- Reciprocity (in directed networks)
- Assortativity degree
- Transitivity
- Mean distance
- Diameter
- Mean degree

These measures help reveal patterns in connectivity and vulnerability, offering insights into how financial systems adapt and destabilize under stress.

Systemic Risk

Following Konstantinov and Fabozzi (2025), the focus in assessing systemic risk lies not solely on its origins—often stemming from external factors, such as macroeconomic shifts, financial distress among key institutions, or sovereign crises—but also on how it spreads and permeates through the network. Thus, systemic risk analysis focuses on the potential failure of individual nodes within the system because their risk exposures can precipitate broader systemic repercussions. Nodes play a pivotal role here because this analysis indicates which nodes are most susceptible to being affected, either simultaneously or in a specific sequence.

Systemic risk is the possibility that a single shock or failure will set off a series of events that could cause the financial system to collapse or be severely disrupted, affecting its functionality and potentially impacting the real economy. It is an emergent characteristic resulting from the interdependence of financial institutions and their group dynamics.

To properly understand the meaning of systemic risk, financial contagion, and spillover, it is necessary to understand the metrics that describe networks and capture interconnectedness. Following Konstantinov and Fabozzi (2025), financial networks are characterized by specific metrics. A thorough understanding of these metrics is critical for understanding network dynamics and interconnectedness. Some common network properties include size, degree, density, and community structure, among others. According to Konstantinov and Fabozzi (2025), "These metrics refer to the overall properties of a network and aim to describe its underlying structure. The structure of network connectedness depends on the underlying information flow process, or how nodes interact over the edges." The metrics that describe the concentration of interconnectedness and the degree of nodal connectedness deserve special attention.

- *Role of density:* Network connectivity and financial stability have a complicated relationship that is frequently characterized as "robust yet fragile." Although fully connected networks offer maximum diversification, moderate connectivity can increase resilience by distributing minor shocks across more institutions. In other words, a network with few links is more fragile than a network with a large number of links. However, higher connectivity has the potential to magnify significant shocks beyond a certain threshold, facilitating quicker contagion and enhancing systemic fragility. As a result, there is a tipping point at which additional links cause the system to become unstable rather than stable.
- *Role of concentration:* Das (2016) showed that even when the overall level of connectivity is the same, financial systems with highly concentrated exposures are typically more susceptible to contagion and systemic risk than systems with more dispersed exposures. In a concentrated system, partners may suffer catastrophic losses if a major counterparty fails.

High-centrality nodes often play a critical role in systemic risk. Hubs—measured by degree, or eigenvector—are systemically important financial institutions whose failure can trigger widespread contagion, making their identification essential for regulation. Bridges, characterized by high betweenness centrality, act as key intermediaries, and their failure can disrupt the flow of capital or information across the network.

Because of both confidentiality and complexity, a major barrier to financial network modeling is limited data. Regulators frequently have access to only a portion of the data, such as aggregate exposures. Although such methods as sparse reconstruction and maximum entropy aid in estimating network structure, they may understate systemic risk if they fail to account for network sparsity or assume uniform link weights.

Financial Contagion and Spillover

Risk propagation in financial networks has been studied using mathematical epidemiology models because of the similarities between financial contagion and disease spread. Financial contagion occurs when a shock or distress event originating in one part of the financial system (e.g., an institution, market segment, or asset class) spreads across the network, potentially causing systemic failures or widespread instability. This process is often compared with epidemic spread or a domino effect, driven by the intricate interconnectedness of financial entities. There are two broad channels of contagion:

- *Direct linkages:* These arise from explicit contractual obligations between financial institutions. For example, if Institution A defaults on a loan owed to Institution B (as in interbank lending), Institution B may suffer direct financial losses. These losses can impair its ability to meet obligations, possibly triggering further defaults.
- *Indirect linkages:* These occur without direct contractual ties, operating through market-wide mechanisms, such as liquidity herding, fire sales of assets, information contagion, or common asset exposures.

Indirect and direct channels often interact. For instance, a direct counterparty loss may induce fire sales, which depress asset values and spark funding runs, escalating the contagion. Network models aim to capture and simulate these complex interactions, offering regulators and analysts a way to assess vulnerability propagation paths.

Case Study: Global Financial Crisis in 2008

Financial shocks can be amplified by network structure, as the 2008 crisis showed. Through a highly interconnected system that included repos, interbank lending, and complex derivatives, such as mortgage-backed securities, collateralized debt obligations, and CDSs, what started as losses in the US subprime mortgage market swiftly spread throughout the world. The September 2008 collapse of Lehman Brothers, a major market node, sparked widespread concerns about counterparty risk. This heightened concern resulted in repos and interbank lending being frozen, demonstrating how the collapse of a major institution can cause systemic instability. The risk of direct contagion was also made evident by AIG's near default. Because AIG had sold significant amounts of CDS protection to big banks, its failure would have resulted in massive losses all at once, necessitating a government bailout to stop further collapse. The crisis was made worse by fire sales.

The global financial crisis emphasized how important it is to adopt a regulatory approach that goes beyond the solvency of individual institutions and specifically takes into account both network interconnectedness and systemic risk. Tools for network analysis became crucial for comprehending these weaknesses. The crisis brought to light the system's "robust-yet-fragile" nature: Interconnectedness that could withstand minor shocks turned into a pathway for catastrophic failure when those shocks grew too big.

Investment Management

Network theory provides useful tools and perspectives for investment management practitioners engaged in portfolio construction, market prediction, and pattern recognition, in addition to risk management and regulatory oversight.

Diversification in Portfolio Construction: Modern Portfolio Theory vs. Network Theory

Expected returns, variances, and the pairwise covariance matrix of assets are the main components of traditional portfolio construction, which is dominated by Markowitz's mean–variance optimization (MVO). MVO has real-world drawbacks, however, such as sensitivity to errors in input estimation (particularly expected returns) and potentially unstable or unduly concentrated allocations. Through the explicit incorporation of the richer structure of asset relationships uncovered by network analysis, network-based approaches seek to improve or offer alternatives. A direct relationship exists between the MVO framework and network centrality scores, which is extensively discussed in Zareei (2019), Ciciretti and Pallotta (2024), and Konstantinov and Fabozzi (2025), among others.

For a minimization of portfolio risk, the mathematical algorithm searches for all possible weights w that minimize the portfolio risk captured by the portfolio variance σ^2_{PF}, which is a weighted product of the individual weights and the variance–covariance matrix, $\Sigma = [\sigma_{ij}]$ (see Konstantinov, Fabozzi, and Simonian 2023). That is,

$$\min_{w} \sigma^2_{PF} = w'\Sigma w$$
$$\text{s.t. } x \geq L,$$

where w is a transposed n-dimensional vector of portfolio weights and Σ is the n-dimensional variance–covariance matrix of the portfolio assets. The variable x represents the portfolio return, while L represents the minimum return value that the portfolio must satisfy. The weights that satisfy the optimizations are

$$w_{min} = \frac{1}{1'\Sigma^{-1}1}\Sigma^{-1}1,$$

where Σ^{-1} is the inverse covariance matrix and 1 represents vectors of ones.

Using the notion of the covariance matrix, Σ, that it is a product of the correlation matrix, Ω, and the diagonal matrix, Δ, whose entries are the variances with $\sigma_i = \sqrt{\sigma_{ij}}$, then the relationship is $\Sigma = \Delta\Omega\Delta$ and the weights subject to the correlation matrix are

$$\frac{1}{1'\Sigma^{-1}1}\Sigma^{-1}1\Omega^{-1}.$$

Considering the mean–variance framework, the diversification, or minimum risk given the expected level of return as measured by the portfolio variance and return $E(r)$, is computed as follows:

$$\min_{w} \sigma^2_{PF} = w'\Sigma w \text{ with } R_{PF} = w'E(r).$$

As a result, the weights of the mean–variance framework are a function of the expected portfolio returns and the covariance matrix of the assets, which gauges risk:

$$w_{mv} = \frac{R_{PF}}{E(r)'\Sigma^{-1}E(r)}\Sigma^{-1}E(r),$$

where Σ^{-1} is the inverse covariance matrix, $E(r)$ represents vectors of expected asset returns, and R_{PF} is the portfolio return. Similarly, we can obtain the weights in the mean–variance framework using the correlation matrix.

In contrast, in network theory, diversification in a network context is identified by weaker relationships between nodes in a graph. The shorter the distances between nodes, the greater the potential impact of risk transmission. This limitation, as noted by Peralta and Zareei (2016) and Zareei (2019), has been addressed and relaxed by Konstantinov (2022), whose approach has found broader application in portfolio allocation, as discussed later in this chapter. Fundamentally, the transfer of risk between nodes is central when using networks for portfolio allocation.

It is important to highlight that centrality scores—such as degree, eigenvector, or alpha centralities—are critical here because they help identify how risk might flow and affect nodes. Degree centrality is not the preferred score, however, because it measures only the immediate connectedness of a node, whereas other centrality metrics capture the more nuanced interconnectedness and influence of a node within the network.

When network theory is applied to portfolio construction, several methods can be used to heighten the diversification of asset portfolios, including the following:

- *Network-based asset selection:* These techniques seek to combine related assets to ensure diversified portfolio exposure across various sources of risk and return.
- *Community detection on association, probabilistic, or statistical networks:* To find clusters of co-moving assets, such algorithms as Louvain, Spinglass, leading eigenvector, or fast greedy provide efficient ways to detect community affiliations with differing factor exposures. Hierarchical clustering or *k*-means methods are applied to asset correlation or distance matrices. Instead of relying on static classifications, selecting representative assets from each cluster can improve diversification and better capture changing market structures. This approach does not require a specific network formulation and is widely used in finance.
- *Graph filtering:* The minimum spanning tree (MST) method highlights important relationships and hierarchies by extracting simplified structures from dense correlation matrices. MST connects all assets with the smallest possible total edge distances. Asset selection can be guided by examining these structures (e.g., central versus peripheral nodes). Some research has shown that portfolios constructed from peripheral MST nodes exhibit strong performance. However, a major drawback of using MST to visualize graphs is that important links might be omitted. Therefore, MST should be used with great caution.
- *Network-based weighting schemes:* Network principles can also guide capital distribution among assets after selection.

- *Centrality-based weighting:* The position of an asset in the network is directly used for weighting. There is an inverse relationship between centrality scores and the weights assigned in the minimum-variance framework. Consequently, peripheral assets are favored, potentially lowering contagion risk. This decision, however, may vary depending on market conditions. Expected returns also play an essential role because centrality-based algorithms focus on risk and interconnectedness but do not account for expected return impacts. Other centrality metrics can be applied to better integrate expected returns.
- *Fundamental networks:* These link assets not only by price movements but also by financial fundamentals (such as profits). To enhance diversification, more capital is allocated to assets that are fundamentally distinct from others. Financial research offers examples of leveraging such information with centrality metrics and risk management tools.
- *Hierarchical risk parity of a community structure:* By using clustering to group similar assets, this model distributes capital among groups to achieve a more balanced risk allocation (see Raffinot 2018).
- *Network risk parity:* A variant of hierarchical risk parity, this method uses an asset's degree of network connectivity based on risk. To reduce exposure to systematic risks that might become systemic, assets less central in the network are given more weight.

As the number of assets increases, these network-based techniques aim to achieve more robust diversification and potentially better risk-adjusted performance than traditional methods relying solely on pairwise statistics. They do so by explicitly leveraging the topological structure of asset relationships.

Market Prediction

Network theory offers frameworks and inputs for forecasting market movements and creating innovative trading tactics. Profitable opportunities or information about future price behavior can be found in the structure of connections between entities. Predictive signals can also be derived from network properties, such as reciprocity, edge density, clustering coefficients, and change in degree centralities.[5] Variations in network density or structure may indicate changes in market volatility or regimes. Spillover effects, in which the actions of related entities affect the performance of a target entity in the future, can be captured by network analysis and applying appropriate network metrics, such as network nodal entropy, Ricci curvature, fragility, or criticality indicators.[6]

One area of quantitative finance that is expanding quickly is the fusion of network science and artificial intelligence, specifically machine learning and deep learning. Although more straightforward techniques can still be competitive, deep neural networks (DNNs) frequently outperform conventional models in forecasting stock returns. Adjusting to shifting market conditions is a major challenge in these models; by modifying regularization according to recent performance, however, such strategies as online early stopping are helpful.

---

[5] See Das (2016) and Konstantinov, Chorus, and Rebmann (2020) for several network-based indicators to predict market behavior.

[6] See Konstantinov and Fabozzi (2025) for extensive discussion and estimation of risk indicators in portfolio management.

Graph neural networks (GNNs) are a category of DNNs that are perfect for modeling financial relationships such as stock interdependencies because they can learn from both node features and network structure, modeling the dependencies between entities. A trading GNN, for instance, has been proposed as a way to calculate the impact of assets, dealers, and their relationships on prices (Wu 2025).

By using attention scores to weight neighbors differently, graph attention networks outperform GNNs in identifying the most pertinent connections. This ability helps with stock prediction and portfolio optimization in financial networks where relationships are complex and have different levels of importance.

Conclusion

As this chapter has shown, applying network theory to financial markets and institutions offers finance professionals useful tools and insightful perspectives. The intricate reality of interconnectedness is not captured by the conventional perspective of isolated actors or homogeneous systems. Network theory provides a strong framework for explicitly modeling these connections.

The integration of AI, the use of multiplex and multilayer networks, developments in explainable AI (XAI) for improved interpretability, and the inclusion of alternative data sources are some new trends in the application of network theory to finance.

Although network analysis provides insightful information, some obstacles exist for applying it in finance. Because data on exposures such as loans or derivatives are often confidential, practitioners are forced to work with partial or proxy data, which introduces potential biases. Complexity, speed, and interpretability must all be balanced in advanced models, which can be computationally taxing, particularly for dynamic or AI-driven models. Additionally, there is no one-size-fits-all tool. The data, research question, and objectives must all be considered when selecting a network structure, metrics, or algorithms. Drawing reliable conclusions requires an understanding of each tool's limitations.

References

Barabási, Albert-László. 2016. *Network Science*. Cambridge, UK: Cambridge University Press.

Bech, Morten L., and Enghin Atalay. 2010. "The Topology of the Federal Funds Market." *Physica A: Statistical Mechanics and Its Applications* 389 (22): 5223–46. doi:10.1016/j.physa.2010.05.058.

Borgatti, Stephen P., Martin G. Everett, Jeffrey C. Johnson, and Filip Agneessens. 2022. *Analyzing Social Networks Using R*. Los Angeles: SAGE Publications.

Ciciretti, Vito, and Alberto Pallotta. 2024. "Network Risk Parity: Graph Theory-Based Portfolio Construction." *Journal of Asset Management* 25 (2): 136–46. doi:10.1057/s41260-023-00347-8.

Clauset, Aaron, Mark E. J. Newman, and Cristopher Moore. 2004. "Finding Community Structure in Very Large Networks." *Physical Review E* 70 (6): 6–11. doi:10.1103/PhysRevE.70.066111.

Das, Sanjiv Ranjan. 2016. "Matrix Metrics: Network-Based Systemic Risk Scoring." *Journal of Alternative Investments* 18 (4): 33–51. doi:10.3905/jai.2016.18.4.033.

Diebold, Francis X., and Kamil Yilmaz. 2015. *Financial and Macroeconomic Connectedness: A Network Approach to Measurement and Monitoring*. New York: Oxford University Press. doi:10.1093/acprof:oso/9780199338290.001.0001.

Konstantinov, Gueorgui S. 2022. "Emerging Market Bonds: Expected Returns and Currency Impact." *Journal of Portfolio Management* 48 (8): 139–58. doi:10.3905/jpm.2022.1.385.

Konstantinov, Gueorgui S., Irene Aldridge, and Hossein B. Kazemi. 2023. "Financial Networks and Portfolio Management." *Journal of Portfolio Management* 49 (9): 190–216. doi:10.3905/jpm.2023.1.525.

Konstantinov, Gueorgui, Andreas Chorus, and Jonas Rebmann. 2020. "A Network and Machine Learning Approach to Factor, Asset, and Blended Allocation." *Journal of Portfolio Management* 46 (6): 54–71. doi:10.3905/jpm.2020.1.147.

Konstantinov, Gueorgui S., and Frank J. Fabozzi. 2025. *Network Models in Finance: Expanding the Tools for Portfolio and Risk Management*. Hoboken, NJ: John Wiley & Sons.

Konstantinov, Gueorgui S., Frank J. Fabozzi, and Joseph Simonian. 2023. *Quantitative Global Bond Portfolio Management*. Singapore: World Scientific Publishing. doi:10.1142/13313.

Newman, Mark. 2010. *Networks: An Introduction*, 1st ed. Oxford, UK: Oxford University Press. doi:10.1093/acprof:oso/9780199206650.001.0001.

Peralta, Gustavo, and Abalfazl Zareei. 2016. "A Network Approach to Portfolio Selection." *Journal of Empirical Finance* 38 (Part A, September): 157–80. doi:10.1016/j.jempfin.2016.06.003.

Raffinot, Thomas. 2018. "Hierarchical Clustering-Based Asset Allocation." *Journal of Portfolio Management* 44 (2): 89–99. doi:10.3905/jpm.2018.44.2.089.

Wasserman, Stanley, and Katherine Faust. 1994. *Social Network Analysis: Methods and Applications*. New York: Cambridge University Press. doi:10.1017/CBO9780511815478.

Wu, Xian. 2025. "Trading Graph Neural Network." Working paper (10 April). doi:10.48550/arXiv.2504.07923.

Yang, Zhao, René Algesheimer, and Claudio J. Tessone. 2016. "A Comparative Analysis of Community Detection Algorithms on Artificial Networks." *Scientific Reports* 6. doi:10.1038/srep30750.

Zareei, Abalfazl. 2019. "Network Origins of Portfolio Risk." *Journal of Banking & Finance* 109 (December). doi:10.1016/j.jbankfin.2019.105663.